UNITED STATES PATENT APPLICATION

FOR

A Cross-Connect with Shared Storage

Inventors:

Peter Yongchun Liu

Prepared by:

Blakely, Sokoloff, Taylor & Zafman LLP
12400 Wilshire Blvd., Suite 700
Los Angeles, California   90025
(310)207-3800

CROSS-REFERENCE TO RELATED APPLICATIONS

Not Applicable.

5

BACKGROUND OF THE INVENTION

Field of the Invention

The invention relates to the field of digital circuitry. More specifically, the invention relates to a cross-connect.

10

Background of the Invention

The current explosion of data traffic traversing the public network is forcing the carriers to change their approach when designing their networks. Time Division Multiplexing (TDM) networks have a place in the telecommunications network, as

15 certain customers require the reservation of portions of bandwidth within the network. The telecommunication industry has sought to maximize the data throughput by developing techniques allowing for the simultaneous transmission of multiple channels on a single telecommunications medium. Data multiplexing is not new to the industry, but introducing new methods in their design are certainly welcomed.

20 A typical TDM network element includes a number of line cards. Data enters and exits the network element through these line cards. The line cards are coupled to a cross-connect that allows for traffic to be selectively transmitted from any of the line cards to any of the other line cards. Thus, the TDM traffic is aligned and switched through the cross-connect.

A typical cross-connect includes three modules: 1) an input Time Slot Interchange (TSI); 2) a switch; 3) and an output TSI. The traffic of each of the line cards is provided to the input TSI. The input TSI includes a buffer for each line card that can reorder the incoming data into a different sequence and/or selectively store

5   parts of it. The output TSI also includes a buffer for each line card. The switch is controlled to selectively provide the output of any one of the buffers in the input TSI to be provided to any one of the buffers of the output TSI. Similar to the buffers of the input TSI, the buffers of the output TSI can selectively reorder and/or selectively store the data. Thus, the buffer of the output TSI can intermix the data from different line

10   cards when the switch and output TSI are so controlled.

A limitation of this prior art method is that it necessitates relatively larger amounts of memory of a type that is typically operated at lower frequencies.

BRIEF SUMMARY OF THE INVENTION

A method and apparatus for sharing storage in a cross-connect is described. According to one embodiment of the invention, a cross-connect includes a number of sets of data input lines and a number of matrices. Each of the sets of data input lines is to be coupled to a different line card. Each of the matrices is coupled to every one of the sets of data input lines. In addition, each of the matrices has a set of data output lines, where the set of data input lines of each of the matrices is to be coupled to a different one of the line cards.

BRIEF DESCRIPTION OF THE DRAWINGS

The invention may best be understood by referring to the following description and accompanying drawings that are used to illustrate embodiments of the invention. In the drawings:

Figure 1 is a diagram of a cross-connect according to one embodiment of the invention.

Figure 2 is a more detailed diagram of one of the matrices 105 from figure 1 according to one embodiment of the invention.

Figure 3 is a diagram representing a single column 205I of matrix entries within the matrix 105J of Figure 2 according to one embodiment of the invention.

DETAILED DESCRIPTION OF THE INVENTION

In the following description, numerous specific details are set forth to provide a thorough understanding of the invention. However, it is understood that the invention may be practiced without these specific details. In other instances, well-known circuits,

5   structures and techniques have not been shown in detail in order not to obscure the invention.

A method and apparatus for sharing storage in a cross-connect is described. While the described cross-connect can be used for many applications, one exemplary use of the cross-connect is in performing Digital Cross-connections and Switching

10   techniques in telecommunication networks. The environment in which the invention is used can encompass many different telecommunication networks where multiplexing methods are applied. In one embodiment, as described in this application, this cross-connect system can be implemented for use in an environment involving Time Division Multiplexing. The implementation of this system has shown a valuable decrease in

15   memory storage necessary to perform digital cross-connections when comparing to prior methods. In addition, the type of memory that can be used to implement the invention (e.g., registers) can be operated at higher frequencies than the type of memory used in prior art cross-connects; thereby allowing for greater density on a chip implementing the cross-connect.

20   Figure 1 is a diagram of a cross-connect according to one embodiment of the invention. In this Figure, the cross-connect 120 is shown receiving traffic from each of a set of line cards 100A through 100X in a network element. It should be understood that the line cards (LCs) are shown in succession where 100A corresponds to a first line card and 100X corresponds to a last line card in the network element; where X is the

number of line cards in that network element coupled to the cross-connect. Figure 1 also shows a set of buffers 110A through 110X corresponding to the set of line cards 100A through 100X. Thus, traffic from line card 100A enters buffer 110A, as so on. Buffers 110 align the incoming traffic. In one embodiment, each of the buffers 110 operates as a FIFO.

Each of the buffers 110A through 110X outputs K bits at a time over a set of data input lines. The value of K will be described in more detail later herein. Figure 1 also shows a set of matrices 105A through 105X corresponding to the set of line cards 100A through 100X. It should be understood that there is a matrix for each of the line cards 100A through 100X. Each of the matrices 105A through 105X provides the output for the corresponding line card 100A through 100X and all function alike. Thus, the output of matrix 105A is provided to the line card 100A, and so on. Each of the matrices is coupled to receive the K input data lines from every one of the buffers 110A through 110X.

Figure 2 is a more detailed diagram of one of the matrices 105 from Figure 1 according to one embodiment of the invention, wherein J denotes one of the matrices 105A through 105X. Thus, the matrix 105J is the matrix to provide the output for the line card 100J. Matrix 105J is made up of rows 200A through 200Q and columns 205A through 205K; where Q represents the number of rows and K is the number of columns in the matrix. The values of Q and K change according to the cross-connect design, and are driven at least in part by the bandwidth requirements. The low value of K would be the minimum size of aligned data (e.g., for a digital cross-connect in a SONET/TDM network element, this value would be 8 because this traffic is byte aligned), and can be increased by multiples of two thereafter to a maximum (e.g., for

SONET/TDM of OC-N, Q*K=N*8 wherein N is greater than 3 and can be increased by multiples of four).

Figure 2 shows a control unit 220 that provides Q sets of control lines. Specifically, Figure 2 shows that each row of the matrix 105J receives a different one
5   the Q sets of control lines. Each set of control lines includes Y+2 lines. The value for Y will described in more detail later herein.

Figure 2 also shows each of the buffers 110A to 110X providing their set of K data lines to the matrix 105J. In particular, Figure 2 shows that the data lines in each set of K data lines have an order according to bit position (positions 1 to K) based on
10   the order of the data being provided by the line cards. In addition, the same bit position input data lines from each of the buffers 110A to 110X are independently grouped to form LC inputs 230A through LC inputs 230K. Thus, the LC inputs 230A includes the data lines for bit position 1 from each of the buffers 110A to 110X, and so on. Therefore, each of the LC inputs 230A through 230K includes X data lines, where X
15   represents the number of line cards 100A through 100X.

For every intersection of rows and columns in the matrix 105J, there exists a matrix entry 235. The LC inputs 230A through 230K are respectively provided to the columns 205A through 205K of the matrix 105J. More particularly, the LC inputs provided to a given column are provided to every matrix entry in that column. In
20   addition, the set of control lines from the control unit 220 provided to a given row are provided to every matrix entry in that row.

Figure 3 is a diagram representing a single column 205I of matrix entries within the matrix 105J of Figure 2 according to one embodiment of the invention, where I denotes one of the columns 205A through 205K. This Figure provides for more detail
25   in describing the matrix entries within a single column of a given matrix 105J. The

coupling of each matrix entry in column 205I to the same LC inputs is illustrated by the coupling of LC inputs 230I to: 1) a matrix entry 305 of row 200A and column 205I; 2) a matrix entry 355 of row 200B and column 205I; and 3) a matrix entry 360 of row 200Q and column 205I.

5   As shown in Figure 3, the matrix entry 305 of row 200A column 205I is conceptually broken down into two major components: a space mux 310 and a TSI element 315. The LC inputs 230I are coupled to the space mux 310 as data inputs. The space mux 310 is also coupled to a set of control lines originating from the control unit 220, as illustrated in Figure 2. While the same LC inputs are coupled as data inputs to

10   the space mux of each matrix entry in the column 205I, a different set of the control lines from the control unit 220 is coupled to each of the space muxes in the column 205I.

As shown in Figure 2, there are Y+2 control lines entering each row of the matrix 105J. Y of these Y+2 control lines are provided to each of the space muxes,

15   where Y is the number of lines required to select one of the LC inputs provided to a matrix entry (thus, $2^Y=K$). Since the same Y control lines are provided to every space mux in a row, a given row is operated to select the data from one of the line cards 100A to 100X at a time. The other 2 control of the Y+2 lines are coupled to other elements in the matrix entry as described later herein.

20   Returning to Figure 3, the TSI element 315 is coupled to the output of the space mux 310 and includes two components: a storage element 320 and a shift element 325. The storage element 320 temporarily stores bit data by employing a timing mux 330 and a store register 335. The control input of the timing mux 330 is coupled to one of the control lines in the set of control lines from control unit 220 for row A. The data

25   input of the timing mux 330 is the output of the space mux 310 and a feedback from the

store register 335. The input to the store register 335 is the output of the timing mux 330.

The shift element 325 is coupled to the output of the storage element 320 and includes a shift mux 340 and a shift register 345. The data inputs of the shift mux 340 are coupled to the output of the store register 345 and zero. The control input of the shift mux 340 is coupled to the last of the Y+2 control lines in the set of control lines from control unit 220 for row A.

The other matrix entries in column 205I are basically the same as the matrix entry 305, and thus only the difference are described. In particular, as previously described a different set of control lines is provided to each row of the matrix, and thus, the set of control lines provided to each of the matrix entries in column 205I is different. In addition, rather than one of the data inputs of the shift mux of the remaining matrix entries in column 205I being zero, that data input is coupled to the output of the preceding matrix entry (e.g., a data input of the shift mux of matrix entry 355 is coupled to the output of the shift register 345).

It should be understood that each matrix entry is reproduced and coupled to its preceding matrix entry; as shown in Figure 3, starting at the first row of the matrix column until reaching the last row of the same matrix column. However, at matrix entry 360 of row 200Q and column 205I, data leaves the shift register 345 of that matrix entry and is outputted to its corresponding line card, in this case line card 100J. Thus, the matrix entries of a column are coupled in series.

The provision of the same set of control lines on a row-by-row basis to every matrix entry on that row, and the provision of the same bit position data input line from each of the line cards on a column-by-column basis to every matrix entry on that column, allows the matrix to, on any given interval, selectively store in any given row

of the matrix the data from one of the line cards. This feature allows the data from a given one of the line cards to be reordered and/or the data from different line cards to be intermixed. For example, a first, second, third and fourth data from the line card 100A could be selectively stored in rows D, A, C, and B respectively. As another

5   example, a first and second data from the line card 100A and a first and second data from the line card 100B could be selectively stored in rows D, A, C, and B respectively.

The interconnection of the matrix entries of a column and the control on a row-by-row basis allows the data in any given row of the matrix to be shifted to the next row of the matrix (or if at the last row, out to the line card). Since the output of each

10   matrix is for a different line card, the matrixes are controlled independently to provide the required output to the different line cards.

Thus, the cross-connect 120 of figure 1 uses as its storage the store registers of each matrix entry of each matrix. Whereas the prior are had separate TSI and switch units with their own storage to reorder, intermix, and store, the storage of the cross-

15   connect 120 is shared for these various operations. Thus, the cross-connect 120 requires less storage than the prior art cross-connect. In addition, since the cross-connect 120 can be implemented with register type storage and since register type storage can operate at higher frequencies, the cross-connect can be implemented at a greater density.

20   The techniques shown in the figures can be implemented using code and data stored and executed on computers. Such computers store and communicate (internally and with other computers over a network) code and data using machine-readable media, such as magnetic disks; optical disks; random access memory; read only memory; flash memory devices; electrical, optical, acoustical or other forms of

25 , propagated signals (e.g., carrier wave, infrared signals, digital signals, etc.); etc. Of

course, one or more parts of the invention may be implemented using any combination of software, firmware, and/or hardware.

While the invention has been described in terms of several embodiments, those skilled in the art will recognize that the invention is not limited to the embodiments

5 described. The method and apparatus of the invention can be practiced with modification and alteration within the spirit and scope of the appended claims. The description is thus to be regarded as illustrative instead of limiting on the invention.